arXiv:1610.00307v1 [cs.CV] 2 Oct 2016

Plug-and-Play CNN for Crowd Motion Analysis: An Application in Abnormal Event Detection

Mahdyar Ravanbakhsh* mahdyar.r@gmail.com
Moin Nabi[2] moin.nabi@unitn.it
Hossein Mousavi[1] hossein.mousavi@iit.it
Enver Sangineto[2] enver.sangineto@unitn.it
Nicu Sebe[2] niculae.sebe@unitn.it
[1] Istituto Italiano di Tecnologia
[2] DISI, University of Trento

Abstract

Most of the crowd abnormal event detection methods rely on complex hand-crafted features to represent the crowd motion and appearance. Convolutional Neural Networks (CNN) have shown to be a powerful tool with excellent representational capacities, which can leverage the need for hand-crafted features. In this paper, we show that keeping track of the changes in the CNN feature across time can facilitate capturing the local abnormality. We specifically propose a novel measure-based method which allows measuring the local abnormality in a video by combining semantic information (inherited from existing CNN models) with low-level Optical-Flow. One of the advantage of this method is that it can be used without the fine-tuning costs. The proposed method is validated on challenging abnormality detection datasets and the results show the superiority of our method compared to the state-of-the-art methods.

1. Introduction

Crowd analysis gained popularity in the recent years in both academic and industrial communities. This growing trend is also due to the increase of population growth rate and the need of more precise public monitoring systems. In the last few years, the computer vision community has pushed on crowd behavior analysis and has made a lot of progress in crowd abnormality detection [10, 27, 8, 3, 17, 16, 13, 11]. Most of these methods mainly rely on complex hand-crafted features to represent the crowd motion and appearance. However, the adoption of hand-crafted features is a clear limitation, as it implies enforcing some a priori knowledge which, in case of complex video surveillance scene, is very difficult to define. Recently, Deep Neural Networks have resurfaced as a powerful tool for learning from big data (e.g., ImageNet [21] with 1.2M images), providing models with excellent representational capacities. Specifically, Convolutional Neural Networks (CNNs) have been trained via backpropagation through several layers of convolutional filters. It has been shown that such models are not only able to achieve state-of-the-art performance for the visual recognition tasks in which they were trained, but the learned representation can be readily applied to other relevant tasks [20]. These models perform extremely well in domains with large amounts of training data. With limited training data, however, they are prone overfitting. This limitation arises often in the abnormal event detection task where scarcity of real-world training examples is a major constraint. Besides the insufficiency of data, the lack of a clear definition of abnormality (*i.e.*, the context-dependent nature of the abnormality) induces subjectivity in annotations. Previous work highlighted the fact that the unsupervised measure-based methods may outperform supervised methods, due to the subjective nature of annotations as well as the small size of training data [26, 23, 25, 14].

Attracted by the capability of CNN to produce a generic semantic representation, in this paper we investigate how to employ CNN features, trained on large-scale image datasets, to be applied in a crowd dataset with few abnormal event instances. This can alleviate the aforementioned problems of supervised methods for abnormality detection, by leveraging the existing CNN models trained for image classification. Besides, training a CNN with images is much cheaper than with videos; therefore, representing a video by means of features learned with static images represents a major saving of computational cost.

The key idea behind our method is to track the changes in the CNN feature across time. We showed that even very

*Work done while M. Ravanbakhsh was interning at DISI, University of Trento.

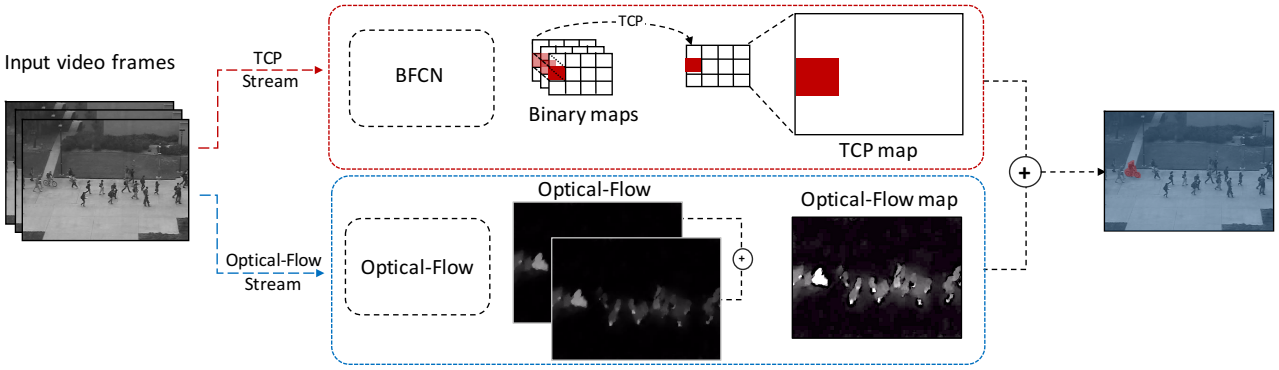

Figure 1. Overview of the proposed method

small consecutive patches may have different CNN features, and this difference captures important properties of video motion. To capture the temporal change in CNN features, we grouped them into a set of binary codes each representing a binary pattern *(prototype)*. Intuitively, in a video block consecutive frames should have similar binary patterns unless they undergo a significant motion. We introduced a simple yet effective statistical measure which captures the local variations of appearance in a video block. We show that combining this simple measure with traditional Optical-Flow, provides the complementary information of both appearance and motion patterns.

Previous Works: Our method is different from [10, 12, 9, 1, 5, 22, 8], which focus on learning models on motion and/or appearance features. A key difference compared to these methods is that they employ classic hand-crafted features (*e.g.*, Optical-Flow, Tracklets, etc.) to model activity patterns, whereas our method proposes using modern deep architectures for this purpose. The advantages of a deep learning framework for anomalous event detection in crowd have been investigated recently in [27]. Nevertheless, learning such deep structures is computationally expensive. In our work, however, a completely different perspective to abnormality detection is picked out. We specifically propose a simple measure-based method which allows the integration of semantic information (inherited from existing CNN models) with low-level Optical-Flow, with minimum additional training cost. This leads to a more discriminative motion representation while maintaining the method complexity to a manageable level. Similar to our proposal, in certain (partial) aspects, a few recent works are mentioning. Most related to our paper is the work by Mousavi *et al*. [14], which introduced a similar measure to capture the commotion of a crowd motion for the task of abnormality detection. Instead of capturing the local irregularity of the low-level motion features (e.g., Tracklets in [14]) or high-level detectors in [15], we propose to represent the crowd motion exploiting the temporal variations of CNN features. This provides the means to jointly employ appearance and motion. Very recently Ravanbakhsh *et al*. [19] proposed a complex feature structure on top of CNN features which can capture the temporal variation in a video for the task of activity recognition. However, to our knowledge this is the first work proposing to employ the existing CNN models for motion representation in crowd.

Method Overview: In general our proposed method includes three steps: 1) Extract CNN-based binary maps from a sequence of input video frames, 2) Compute a measure upon the extracted CNN-binary maps to discover motion patterns 3) Fuse with low-level motion feature (Optical-Flow) to find the refined motion segments. For a given video, first the video frames are input to a Fully Convolutional Network (FCN) sequentially. Then we propose a binary layer to quantize the high-dimensional feature maps into compact binary patterns. The binary quantization layer is a convolutional layer in which the weights are initialized with an external hashing method. The binary layer produces binary patterns for each patch corresponding to the receptive field of the FCN, called *binary map*. The output binary maps preserve the spacial relations in the original frame, which is useful for localization tasks. Then, a histogram is computed over output binary maps for aggregating binary patterns of spatio-temporal blocks. In the next step, a *irregularity* measure is computed over these histograms. Eventually, all the computed measures over all video blocks are concatenated, up-sampled to the original frame size, and fused with Optical-Flow to localize the abnormality. In the rest of this paper we describe each part in detail.

Contributions: Our major contributions in this paper are: *(i)* We introduce a novel Binary Quantization Layer, *(ii)* We propose a Temporal CNN Pattern measure to represent motion in crowd, *(iii)* The proposed method is validated on challenging abnormality detection datasets and the results show the superiority of our method compared to the state-of-the-art methods.

The rest of the paper is organized as follows: the idea of

our Binary FCN is introduced in Sec. 2. In Sec. 3 we explain the proposed measure, and the fusion of the features explained in Sec. 4. The experiments regarding our introduced approaches and a discussion on the obtained results is presented in Sec. 5.

2. Sequential Binary Fully Convolutional Network (BFCN)

In this section, we present the sequential Fully Convolutional Network (FCN) which creates the binary maps for video frames. The proposed architecture contains two main modules: *1) The Convolutional feature maps*, and *2) Binary map representation of local features*. In the following, we describe each part in details.

Frame-based Fully Convolutional Network Structure: Early layers of convolutions in deep nets present local information about the image, while deeper convolutional layers contain more global information. Fully connected layers represent high-level information and usually can be used for classification and recognition tasks. It has been shown that deep net models which are trained on the ImageNet [21] encodes semantic information, thus can address a wide range of recognition problems [20, 2]. Since, FCNs do not contain a fully-connected layer they preserve relative spatial coordinates between the input image and the output feature map. Hence a feature in the output map corresponds to a large receptive field of the input frame. Beside, FCNs are capable to accept input images of different sizes and return feature maps of different sizes as output. In light of the above, this deep net not only provides an appropriate structure to extract local and global information about the image, but also preserves spatial relations, which is a big advantage for a localization task.

Convolutional Feature maps: The image intensity values are not representing any global information about the image. To tackle the gap between intensity and high-level information we choose the output of the last convolutional layer to extract feature maps. These components provide global information about the objects in the scene. To extract convolutional feature maps, we used a pre-trained AlexNet [6] model. Original AlexNet contains 5 convolutional layers and two fully connected layers. In order to obtain spatially localizable feature maps, we feed the output feature maps of the last convolutional layer into our binary quantization layer. Figure 2 illustrates the layout of our network.

Binary Quantization Layer: In order to generate a joint model for image segments there is a need to cluster feature components. Clustering the extracted high-dimensional feature maps comes with a high computational cost. The other problem with clustering is the need to know a priori the number of cluster centres. One possible approach to avoid extensive computational costs and to obtain reasonable efficiency is clustering high-dimensional features with a hashing technique for generating small binary codes for each feature vector. A 24-bits binary code can address 2^{24} cluster centres, which is impossible to be handled by the clustering methods. Also, this binary map can be simply represented as a 3-channels RGB image. Dealing with binary codes comes with a lower computational cost and a higher efficiency in comparison to other clustering methods. The other advantage of using a hashing technique in comparison with clustering is the capability of embedding the pre-trained hash as a layer inside the network.

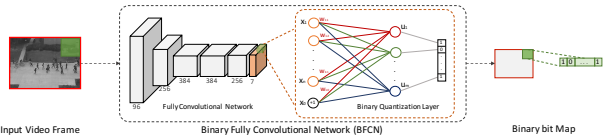

Figure 2. Fully Convolutional Net with Binary Quantization Layer

Binary Encoding Layer: Encoding feature maps to binary codes is done by Iterative Quantization Hashing (ITQ)[4], which is a hashing method for unsupervised-learning binary codes. Training the ITQ is the only training cost in the proposed method, which can be done only once on a subset of data.The ITQ projects each high-dimensional feature vector into a binary space. We use the hashing weights, which are learned by ITQ, to build up a binary encoding module as a last layer (denoted by $hconv6$) in the network architecture. Specifically, inspired by [18] we implement this layer as a set of convolutional filters (shown in different colors in Fig 2), followed by a sigmoid activation function. The number of these filters is equal to the size of the binary code and the weights are computed through ITQ. Finally, the binarization step has been done externally by thresholding the output of the sigmoid function.

Specifically, for $X = \{x_1, x_2, ..., x_n\}$ a feature vector of $pool5$, the output of $hconv6$ is defined by $hconv6(X) = XW_i$, where $W_i = \{w_{i1}, w_{i2}, ..., w_{in}\}$ are the weights for the i^{th} neuron. The non-linearity is provided with a sigmoid function $\upsilon = \sigma(hconv6(X))$, and the threshold function is defined by:

$$g(\upsilon) = \begin{cases} 0, & \upsilon \leq 0.5 \\ 1, & \upsilon > 0.5 \end{cases} \quad (1)$$

Eventually, for any given video frame our network returns a spatial-aware binary bitmap, which can be used for localization tasks.

Such a binary quantization layer can be plugged into the net as a pre-trained module, with the possibility of fine-tuning with back-propagation in an end-to-end fashion. However, in the abnormality task, due to lack of data there is no possibility for fine-tuning. However, it could be trained for other tasks.

Sequential BFCN: Let $\mathbf{v} = \{f_t\}_{t=1}^{T}$ be an input video,

Method	Ped1 (frame level)		Ped1 (pixel level)		Ped2 (frame level)	
	ERR	AUC	ERR	AUC	ERR	AUC
MPPCA [5]	40%	59.0%	81%	20.5%	30%	69.3%
Social force(SF) [10]	31%	67.5%	79%	19.7%	42%	55.6%
SF+MPPCA [9]	32%	68.8%	71%	21.3%	36%	61.3%
SR [1]	19%	—	54%	45.3%	—	—
MDT [9]	25%	81.8%	58%	44.1%	25%	82.9%
LSA [22]	16%	92.7%	—	—	—	—
Detection at 150fps [8]	15%	91.8%	43%	63.8%	—	—
AMDN (early fusion) [27]	22%	84.9%	47.1%	57.8%	24%	81.5%
AMDN (late fusion) [27]	18%	89.1%	43.6%	62.1%	19%	87.3%
AMDN (double fusion) [27]	16%	92.1%	40.1%	67.2%	17%	90.8%
Proposed Method	8%	95.7%	40.8%	64.5%	18%	88.4%

Table 1. Comparison with state-of-the-art on UCSD dataset: reported ERR (Equal Error Rate) and AUC (Area Under Curve). The values of previous methods are reported from [27].

where f_t is the t-th frame of the video, and T is the number of frames. The frames of the video are fed into the network sequentially. The output for a single frame $f_t \in \mathbf{v}$, is an encoded binary bit map (prototype), denoted as m_t. All the binary maps are stacked and delivered as the final output of a given video, *i.e.*, $\mathbf{M} = \{m_t\}_{t=1}^{T}$.

3. Temporal CNN Pattern (TCP)

In this section we describe our proposed method to measure abnormality in a video sequence.

Overlapped Video Blocks: The first step is extracting video blocks from the video clip. As mentioned in Sec. 2, for a given video frame f_t the output of our FCN is a binary bit map denoted by m_t, which is smaller in size than the original image. In the binary bit map representation, each pixel describes a corresponding region in the original frame. This feature map partitions the video frame into a certain number of regions, which we called $patch$ denoted by p_t^i where t is the frame number and i is the i-th patch in the frame. b_t^i is a set of corresponding patches along consecutive frames. The length of a video blocks is fixed, and the middle patch of the video block is selected to indicate the entire video block. If the length of a video block b_t^i is $L+1$, it starts $L/2$ frames before the frame t and ends $L/2$ frames after that, namely $\{b_t^i\} = \{p_l^i\}_{l=t-L/2}^{t+L/2}$. To capture more fine information, the video block b_t^i has n frames overlapped with the next video block b_{t+1}^i.

Histogram of Binary Codes: A video block is represented with a set of binary codes (prototypes), which encode the similarity in appearance. We believe that observing the changes of prototypes during time is a clue to discover motion patterns. Toward this purpose, for each video block b_t^i a histogram h_t^i is computed to represent the distribution of prototypes in the video block.

TCP Measure: Similarly to the commotion measure [14], to obtain the TCP measure for a given video block b_t^i, the *irregularity* of histogram h_t^i is computed. This is done by considering the fact that, if there is no difference in the appearance, then there is no change in descriptor features and consequently there is no change in the prototype representation. When the pattern of binary bits changes, it means different appearances are observed in the video block and this information is used to capture motion patterns. The *irregularity* of the histogram defined as the non-uniformity of the distribution in the video block. A uniform distribution of a histogram shows the presence of several visual patterns in a video block. The higher diversity of the prototypes on a video block leads to the lower *irregularity* of the histogram. More uniform histograms increase the chance of abnormality. Such *irregularity* in appearance along the video blocks either is generated by noise or is the source of an anomaly. We took advantage of this fact to present our TCP measure.

The TCP measure for each video block b_t^i, is computed by summing over the differences between the prototype samples in h_t^i and the dominant prototype. The dominant prototype is defined as the most frequent binary code in the video block, which has the maximum value (mode) in the histogram h_t^i .

Let H^n represent the histogram of binary codes of all patches $\{p_t^i\}$ in the video block b_t^i denoted by $\{H^n\}_{n=1}^{N}$, where N is the number of patches in the video block. The aggregated histogram for block b_t^i is computed as $\mathcal{H}_t^i = \sum_{n=1}^{N} H^n$. The aggregated histogram $\mathcal{H}_t^i$ represents the distribution of appearance binary codes over video block b_t^i, which is used to compute the TCP measure as:

$$tcp(b_t^i) = \sum_{j=1}^{|\mathcal{H}_t^i|} ||\mathcal{H}_t^i(j) - \mathcal{H}_t^i(j_{max})||_2^2 \qquad (2)$$

where $|.|$ is the number of bins of the histogram, $||.||_2$ is the L2-norm, and the dominant appearance index over the

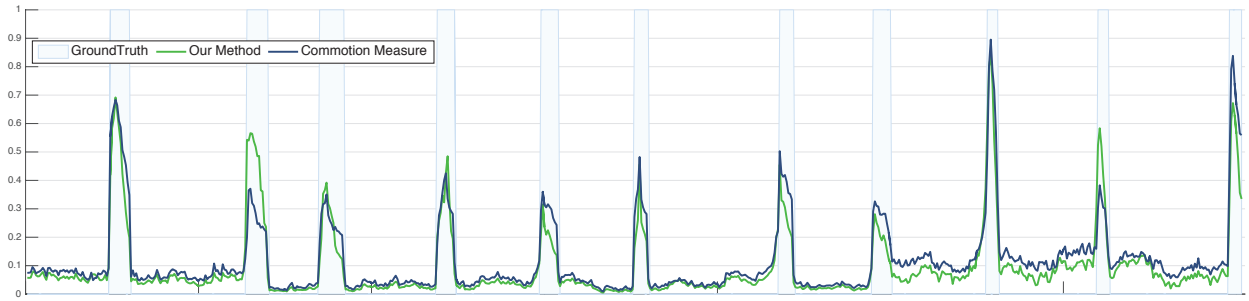

Figure 3. Frame-level anomaly detection results on UMN dataset: Our method compares to Commotion Measure [14]

video block is denoted by j_{max} (*i.e.*, the mode of $\mathcal{H}_t^i$).

TCP Map: To create a spatial map of the TCP measure c_t for any given frame f_t, the TCP measure is computed for all video blocks b_t^i, and we assign the value of c_t^i to a patch that is temporally located at the middle of the selected video block. The output c_t is a map with the same size as the binary map m_t which contains TCP measure values for each patch in the frame f_t. Finally, the TCP maps are extracted for the entire video footage. We denote the TCP map for frame f_t as $c_t = \{c_t^i\}_{i=1}^I$, where I is the number of patches in the frame.

Up-sampling TCP Maps: Since the frame will pass through several convolution and pooling layers in the network, the final TCP map is smaller than the original video frame. To localize the exact region there is a need to produce a map of the same size as the input frame. Toward this purpose, the TCP value in the map is assigned to all pixels in the corresponding patch of the frame on the up-sampled TCP map.

4. Fusion with Optical-Flow Maps

Since the Up-sampled TCP map can only detect the coarse region of abnormality, we propose to fuse Optical-Flow with the TCP maps in order to have a more accurate localization.

Optical-Flow Maps: The Optical-Flow is extracted from each two consecutive frames. However the TCP map is computed for each L frames. To be able to fuse the Optical-Flow with the corresponding extracted TCP map, an aligned Optical-Flow map is constructed. Suppose that f_t and f_{t+1} are two consecutive frames from video $\mathbf{v} = \{f_t\}_{t=1}^T$, Optical-Flow map of_t is the same size of the frames and represents the Optical-Flow values transition between the two frames. Optical-Flow values extracted for entire video footage $\mathbf{v}$ and stacked as Optical-Flow sequences $\{d_t\}_{t=1}^{T-1}$. Finally, similar to the overlapped video block extraction protocol, overlapped Optical-Flow maps are computed. If the length of a video block p_t^i is $L+1$, then the corresponding Optical-Flow map d_t^i is the sum of all Optical-Flow values over the corresponding region i-th as $d_t^i = \sum_{l=1}^{L} d_t^i(l)$.

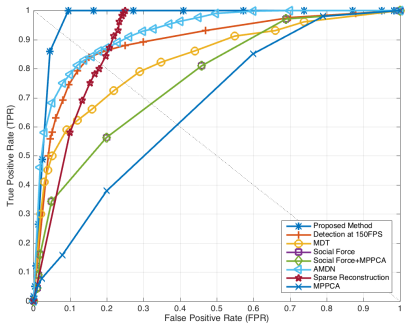

(a) ROC curve of frame level (anomaly detection)

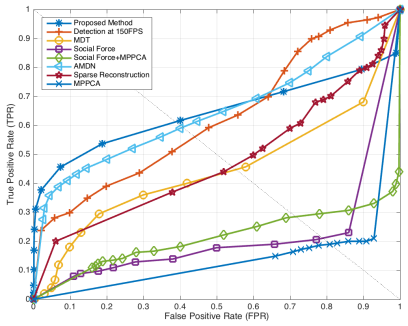

(b) ROC curve of pixel level (anomaly localization)

Figure 4. Frame level and Pixel level comparison ROC curves of Ped1 (UCSD dataset).

The Optical-Flow map for entire frame f_t is described as $d_t = \{d_t^i\}_{i=1}^I$.

Feature Fusion: The extracted Optical-Flow maps and computed TCP maps for each video frame are fused together with importance factors α and β to create motion segment map $mseg_t = \alpha d_t + \beta c_t$, $mseg = \{mseg_t\}_{t=1}^T$, where, $\{mseg\}$ is the extracted motion segments along the entire video $\mathbf{v}$. The importance factors indicates the influence of each fused map in the final segment motion map, we simply select 0.5 for both α and β.

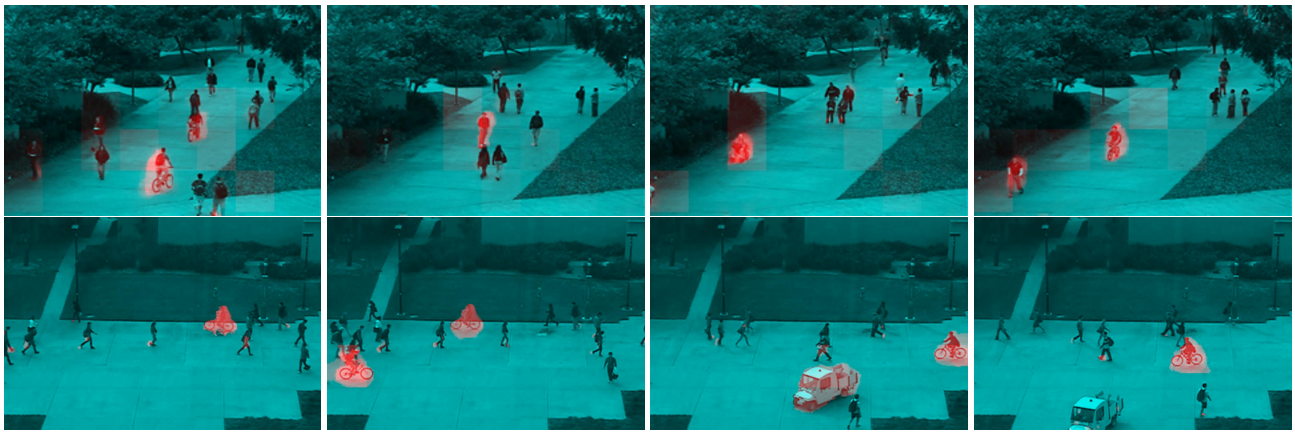

Figure 5. Sample results of anomaly localization on UCSD: Top row selected from Ped1, and bottom row Ped2

5. Experimental Results

In this section, we evaluate our method over two well-known crowd abnormality datasets and compare our results with state of the art. The evaluation has been performed with both a *pixel-level* and a *frame-level* protocol, under standard setup. The rest of this section is dedicated to describing the evaluation datasets, the experimental setup and the reporting the results quantitatively and qualitatively.

Datasets and Experimental Setup: In order to evaluate our method two standard datasets: UCSD Anomaly Detection Dataset [9] and UMN SocialForce [10] . The **UCSD dataset** is split into two subsets *Ped1* and *Ped2*. *Ped1* contains 34/16 training/test sequences with frame resolution 238×158. Video sequences consist of 3,400 abnormal frame samples and 5,500 normal frames. *Ped2* includes 16/12 training/test video samples, with about 1,600 abnormal frames and 350 normal samples. This subset is captured from different scenes than *Ped1*, and the frames resolution is 360×240. This dataset is challenging due to different camera view points, low resolution, different types of moving objects across the scene, presence of one or more anomalies in the frames. The **UMN dataset** contains 11 video sequences in 3 different scenes, and 7700 frames in total. The resolution is 320×240. All sequences start with a normal scene and end with abnormality section.

In our experiments to initialize the weights of $hconv6$ an ITQ is applied on the train set of UCSD pedestrian dataset with a 7-bits binary code representation, which addresses 128 different appearance classes. Video frames are fed to the BFCN sequentially to extract binary bit maps. All video frames are resized to 460×350, then BFCN for any given frame returns a binary bit map with resolution 8×5, which splits the frame into a 40-region grid. The length of video block extracted from a binary map is fixed to $L = 14$ with 13 frames overlapping. The TCP measure is normalized over the entire video block sequence, then a threshold $th <$ 0.1 is applied for detecting and subtracting the background region in the video.

Method	AUC
Optical-Flow [10]	0.84
SFM [10]	0.96
Chaotic Invariants [24]	0.99
Sparse Reconstruction [1]	0.97
Commotion [14]	0.98
TCP (proposed method)	0.98

Table 2. Results on UMN dataset. The values of previous methods are reported from [14].

Optical-Flow feature maps are extracted to fuse with our computed features on the TCP measure maps. The fusion importance factor set to 0.5 equally for both feature sets. These motion segment maps are used to evaluate the performance of our method on detection and localization of anomalous motions during video frames.

5.1. Quantitative Evaluation

Frame Level Anomaly Detection: This experiment aims at evaluating the performance of anomaly detection along the video clip. The criterion to detect a frame as abnormal is based on checking if the frame contains at least one abnormal pixel. To evaluate the performances the detected frame is compared to ground truth frame label regardless of the location of the anomalous event. The procedure is applied over range of thresholds to build the ROC curve.

In the frame level experiments a frame is detected as abnormal if at least it contains one abnormal pixel: in this case the abnormal label is assigned to the whole frame. We compared our method with state-of-the-art methods. We compare our method with state of the art in detection performance on UCSD ped1 and ped2 datasets. The result is shown in Table 1, beside the ROC curves on Figure 4.

The proposed method is also evaluated on UMN dataset. Figure 3 shows the computed TCP for each frame illustrated as a signal (green). We compared TCP with commotion measure (blue). The overall TCP value for a frame is computed from the sum of TCP measures over the patches in a frame.

Pixel Level Anomaly Localization: The goal of the pixel level evaluation is to measure the accuracy of anomalous event localization. Following [7], detected abnormal pixels are compared to pixel level groundtruth. A true positive prediction should cover at least 40% of true abnormal pixels over groundtruth, otherwise counted as a false positive detection.Figure 4 shows the ROC curves of the localization accuracy over USDC Ped1 and Ped2. We compare our method with state of the art in accuracy for localization. Result is presented in Table 1.

Discussion: In our experiments we observed that in most of the cases the proposed method hit the abnormality correctly in terms of detection and localization. Only in case of double fusion in AMDN method [27] our measure achieved lower accuracy in anomaly localization, while the anomaly detection performance always performs better than all the state of the art methods. Note that the proposed method is not taking advantage of any kind of learning in comparison with the others. The proposed method can be effectively exploited to detect and localize anomaly with no additional learning costs. Qualitative results on Ped1 and Ped2 are shown in Figure. 5. The figure shows we could successfully detect different abnormality sources (like cars, bicycles and skateboards) even in the case in which the object can not be recognized by visual appearance alone (*e.g.*, the skateboard). The last column in Figure 5 shows the confusion cases, which not detect the abnormal object (the car) and detect normal as abnormal (the pedestrian).

6. Conclusions

In this paper we addressed the problem of abnormality detection in crowd scenes. We employed a Fully Convolutional Network as a pre-trained model and plugged an effective binary quantization layer as the final layer to the net. Our method provides both spatial consistency as well as low dimensional semantic embedding. We then introduced a simple yet effective unsupervised measure to capture temporal CNN patterns in video frames. We showed that combining this simple measure with traditional Optical-Flow provides us with the complementary information of both appearance and motion patterns. The qualitative and quantitative results on the challenging datasets show that our method outperforms the state-of-the-art methods. As future work, we will study plugging a TCP measure layer and fine-tuning this layer with back-propagation and end-to-end training of abnormality detection. Moreover, exploring of the approach for the task of spatio-temporal motion segmentation would be a potential direction.

References

[1] Y. Cong, J. Yuan, and J. Liu. Sparse reconstruction cost for abnormal event detection. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 3449–3456. IEEE, 2011.

[2] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *arXiv preprint arXiv:1310.1531*, 2013.

[3] R. Emonet, J. Varadarajan, and J.-M. Odobez. Multi-camera open space human activity discovery for anomaly detection. In *Advanced Video and Signal-Based Surveillance (AVSS), 2011 8th IEEE International Conference on*, pages 218–223. IEEE, 2011.

[4] Y. Gong and S. Lazebnik. Iterative quantization: A procrustean approach to learning binary codes. In *CVPR*, 2011.

[5] J. Kim and K. Grauman. Observe locally, infer globally: a space-time mrf for detecting abnormal activities with incremental updates. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 2921–2928. IEEE, 2009.

[6] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.

[7] W. Li, V. Mahadevan, and N. Vasconcelos. Anomaly detection and localization in crowded scenes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 36(1):18–32, 2014.

[8] C. Lu, J. Shi, and J. Jia. Abnormal event detection at 150 fps in matlab. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2720–2727, 2013.

[9] V. Mahadevan, W. Li, V. Bhalodia, and N. Vasconcelos. Anomaly detection in crowded scenes. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 1975–1981, 2010.

[10] R. Mehran, A. Oyama, and M. Shah. Abnormal crowd behavior detection using social force model. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 935–942. IEEE, 2009.

[11] H. Mousavi, H. K. Galoogahi, A. Perina, and V. Murino. Detecting abnormal behavioral patterns in crowd scenarios. In *Toward Robotic Socially Believable Behaving Systems-Volume II*, pages 185–205. Springer International Publishing, 2016.

[12] H. Mousavi, S. Mohammadi, A. Perina, R. Chellali, and V. Murino. Analyzing tracklets for the detection of abnormal crowd behavior. In *Applications of Computer Vision (WACV), 2015 IEEE Winter Conference on*, pages 148–155. IEEE, 2015.

[13] H. Mousavi, M. Nabi, H. K. Galoogahi, A. Perina, and V. Murino. Abnormality detection with improved histogram of oriented tracklets. In *International Conference on Image Analysis and Processing*, pages 722–732. Springer, 2015.

[14] H. Mousavi, M. Nabi, H. Kiani, A. Perina, and V. Murino. Crowd motion monitoring using tracklet-based commotion

measure. In *Image Processing (ICIP), 2015 IEEE International Conference on*, pages 2354–2358. IEEE, 2015.

[15] M. Nabi, A. Bue, and V. Murino. Temporal poselets for collective activity detection and recognition. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 500–507, 2013.

[16] H. Rabiee, J. Haddadnia, H. Mousavi, M. Nabi, V. Murino, and N. Sebe. Emotion-based crowd representation for abnormality detection. *arXiv preprint arXiv:1607.07646*, 2016.

[17] R. Raghavendra, M. Cristani, A. Del Bue, E. Sangineto, and V. Murino. Anomaly detection in crowded scenes: A novel framework based on swarm optimization and social force modeling. In *Modeling, Simulation and Visual Analysis of Crowds*, pages 383–411. Springer New York, 2013.

[18] M. Ravanbakhsh, H. Mousavi, M. Nabi, M. Rastegari, and C. Regazzoni. Cnn-aware binary map for general semantic segmentation. In *Image Processing (ICIP), 2016 IEEE International Conference on*. IEEE, 2016.

[19] M. Ravanbakhsh, H. Mousavi, M. Rastegari, V. Murino, and L. S. Davis. Action recognition with image based cnn features. *arXiv preprint arXiv:1512.03980*, 2015.

[20] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. In *CVPRW*, 2014.

[21] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.

[22] V. Saligrama and Z. Chen. Video anomaly detection based on local statistical aggregates. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2112–2119. IEEE, 2012.

[23] A. A. Sodemann, M. P. Ross, and B. J. Borghetti. A review of anomaly detection in automated surveillance. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 42(6):1257–1272, 2012.

[24] S. Wu, B. E. Moore, and M. Shah. Chaotic invariants of lagrangian particle trajectories for anomaly detection in crowded scenes. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 2054–2060. IEEE, 2010.

[25] T. Xiang and S. Gong. Video behavior profiling for anomaly detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 30(5):893–908, 2008.

[26] G. Xiong, J. Cheng, X. Wu, Y.-L. Chen, Y. Ou, and Y. Xu. An energy model approach to people counting for abnormal crowd behavior detection. *Neurocomputing*, 83:121–135, 2012.

[27] D. Xu, E. Ricci, Y. Yan, J. Song, and N. Sebe. Learning deep representations of appearance and motion for anomalous event detection. *arXiv preprint arXiv:1510.01553*, 2015.